

Ingest Thread Performance

Nick Singer

28 August 1997

Agenda

Test of ECS Performance: Source-to-Ingest-to-Archive Thread



Ingest requirements

High-level architecture

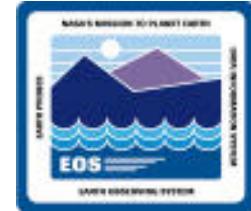
Thread description

Hardware configurations

Results

Conclusions & next steps

Ingest Requirements at Launch



Goddard Space Flight Center (GSFC)

- 70 GB/day for MODIS via EDOS

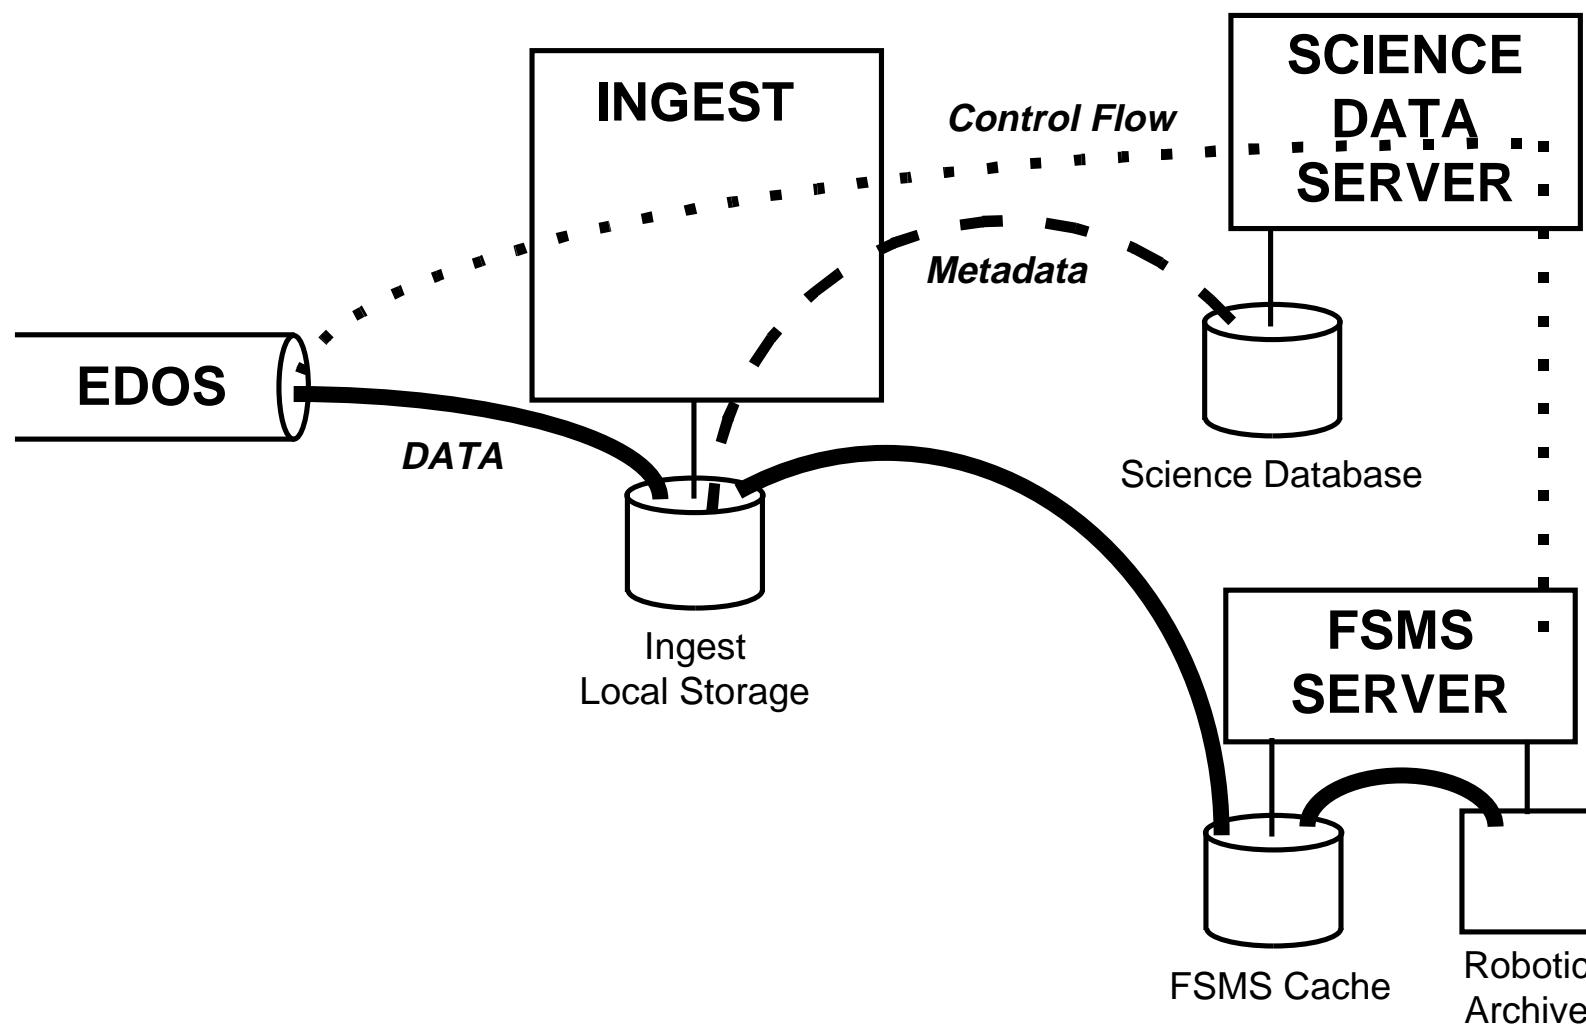
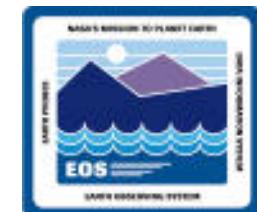
Langley Research Center (LaRC)

- 49 GB/day for MISR, CERES, and MOPPIT via EDOS

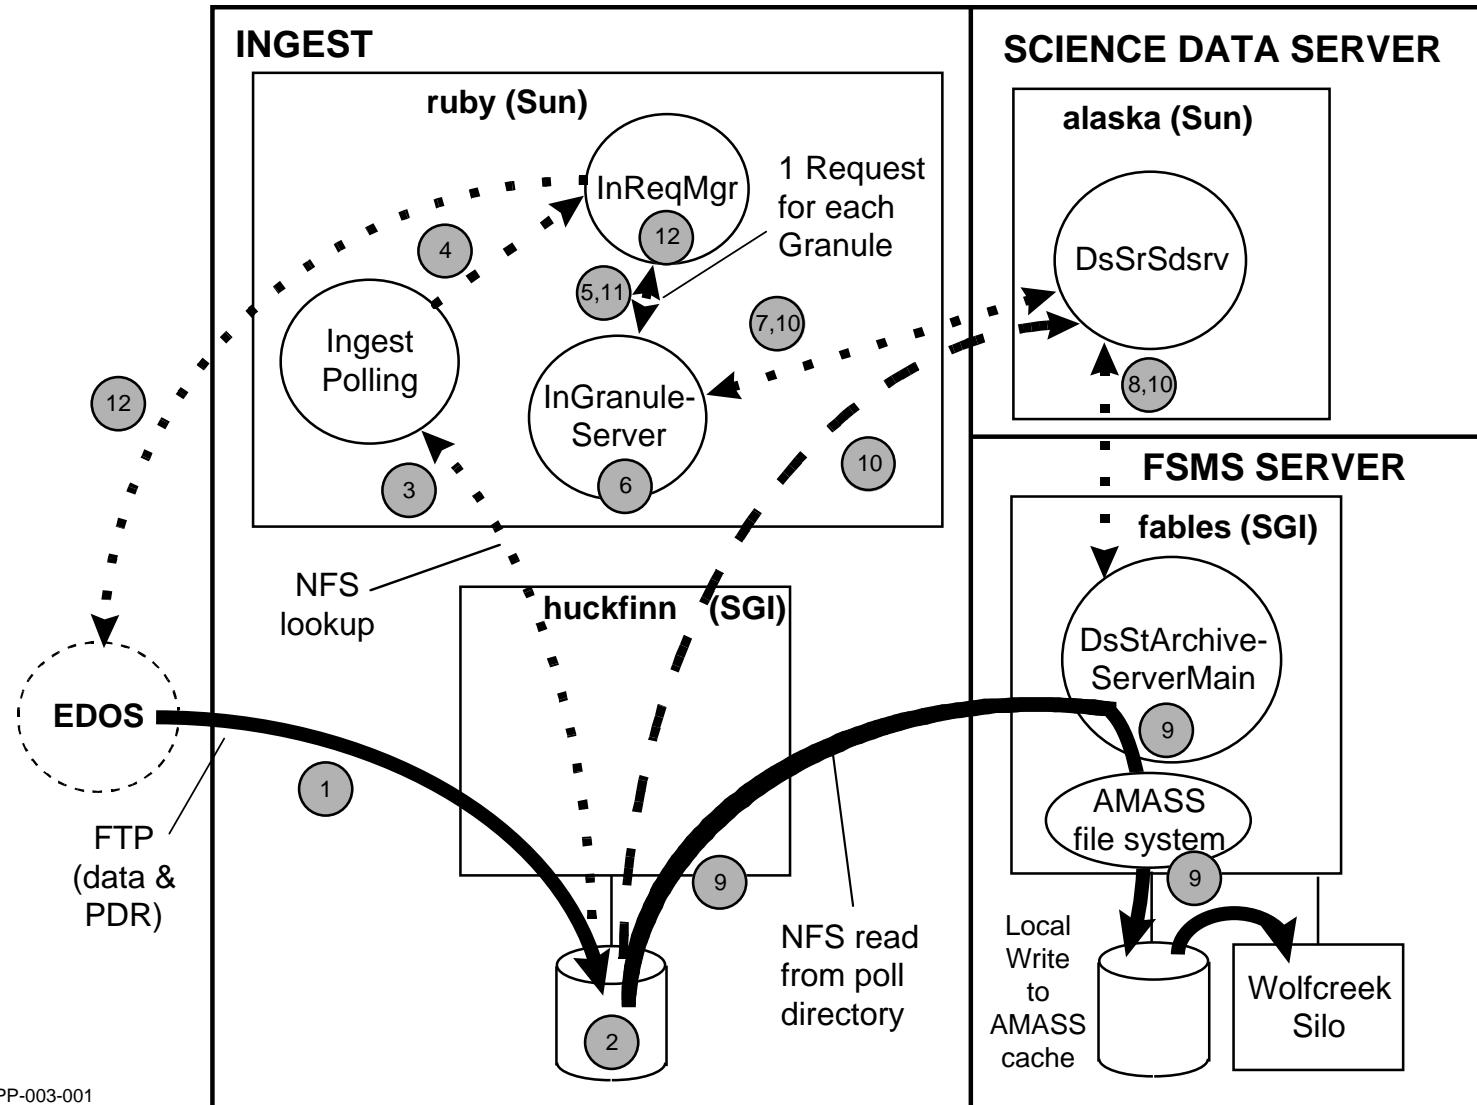
EROS Data Center (EDC)

- 139 GB/day for Landsat-7 via LPS
- 135 GB/day for ASTER from D3 tape

Ingest Thread High-Level Architecture



Source-to-Ingest-to-Archive Thread

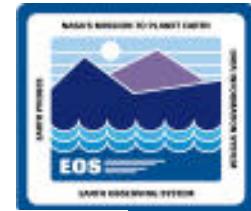


Hardware Configurations

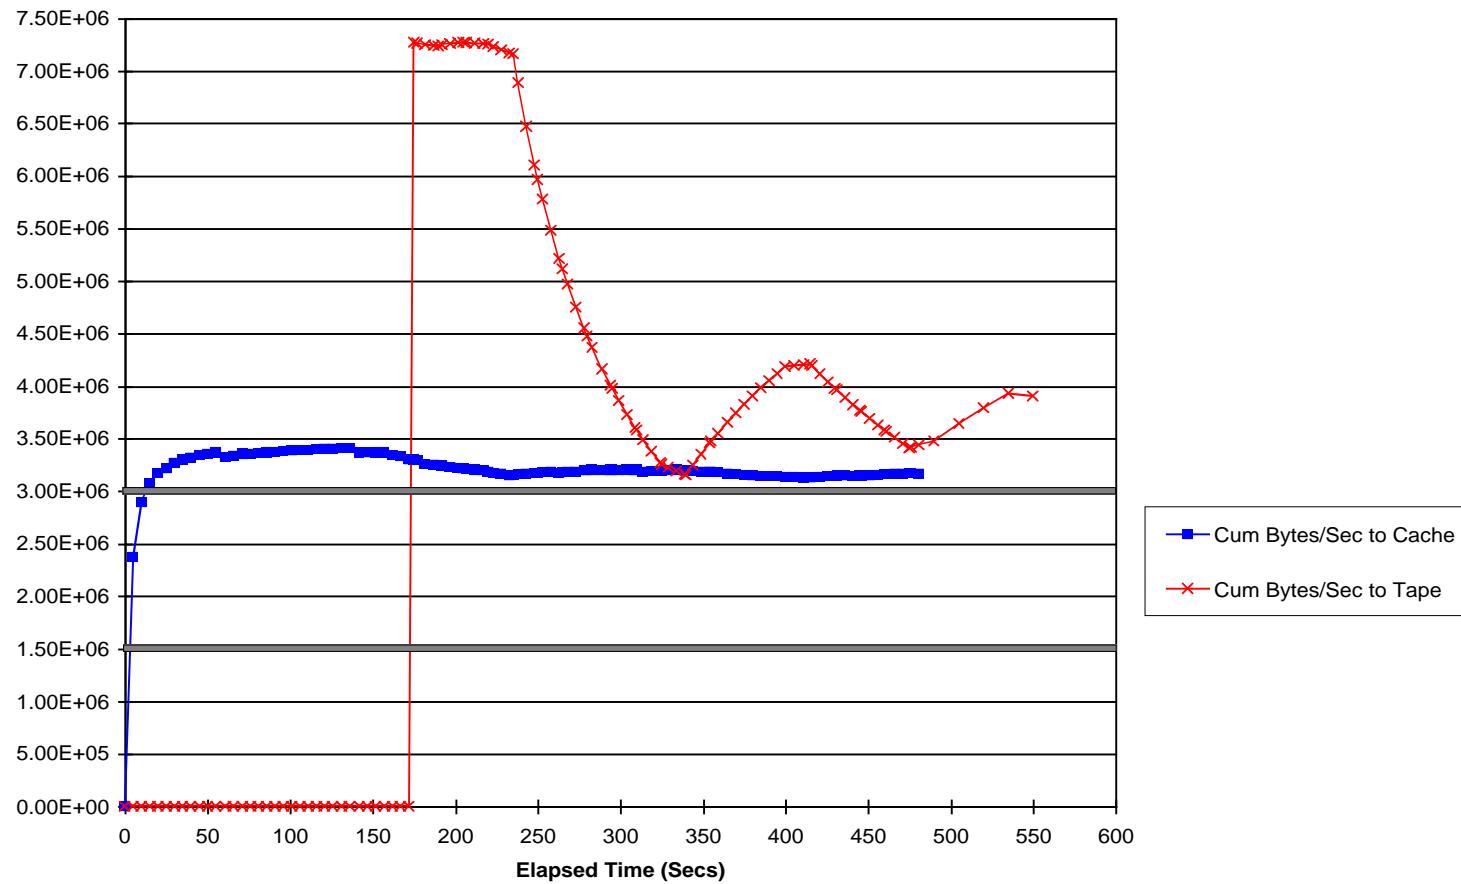


<i>Mini-DAAC</i>	<i>GSFC DAAC</i>
Alaska (Science Data Server) <ul style="list-style-type: none">• Sun UltraSPARC-II, 2 processors @ 248 MHz• 512 MB RAM	g0acs04 (Science Data Server) <ul style="list-style-type: none">• Sun UltraSPARC-II, 2 processors @ 248 MHz• 512 MB RAM
Ruby (Ingest Server) <ul style="list-style-type: none">• Sun UltraSPARC, 2 processors @ 168 MHz• 512 MB RAM	g0icg01 (Ingest Server & Polling Dir.) <ul style="list-style-type: none">• SGI Challenge, 2 R4400 processors @ 200 MHz• 256 MB RAM• 2 RAID controllers
Huckfinn (Ingest Server & Polling Dir.) <ul style="list-style-type: none">• SGI Challenge, 1 R4400 processor @ 150 MHz• 256 MB RAM• 1 RAID controller	
Fables (FSMS Server) <ul style="list-style-type: none">• SGI Challenge, 4 R10000 processors @ 194 MHz• 512 MB RAM• 4 RAID controllers• STK Wolfcreek w/4 D3 drives	g0drg01 (FSMS Server) <ul style="list-style-type: none">• SGI Challenge XL, 8 R10000 processors @ 194 MHz• 512 MB RAM• 8 RAID controllers• STK Powderhorn w/8 D3 drives

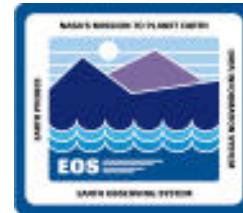
Without ftp Contention—Rates



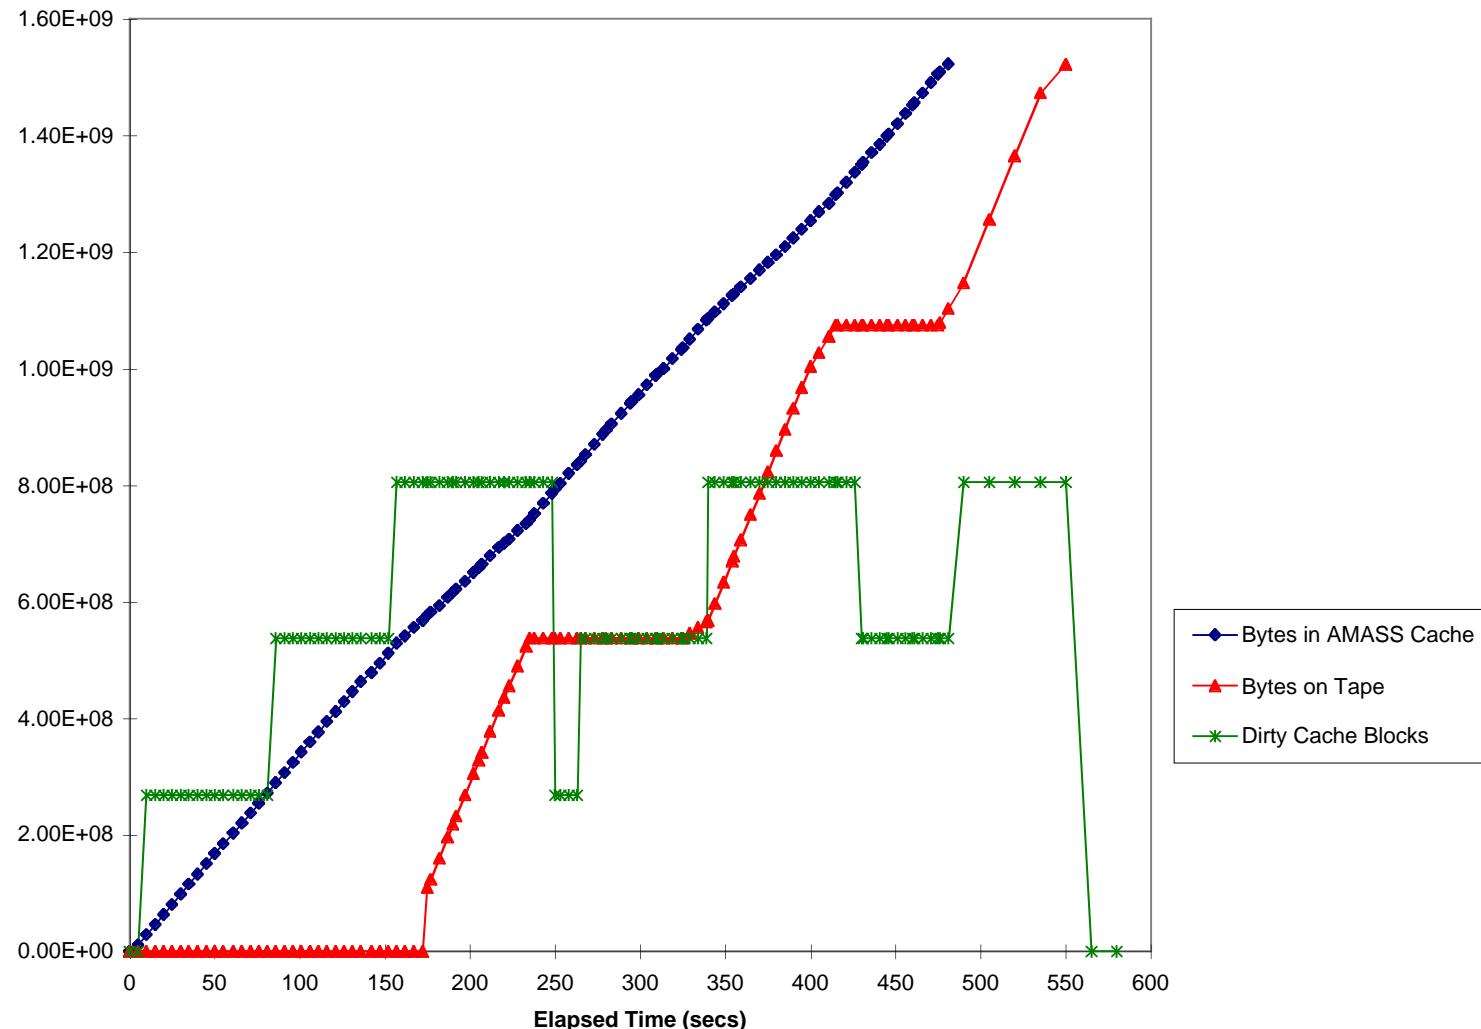
Copy from Ingest to AMASS Cache to AMASS Tape
No ftp Contention at Ingest Source



Without ftp Contention—Details



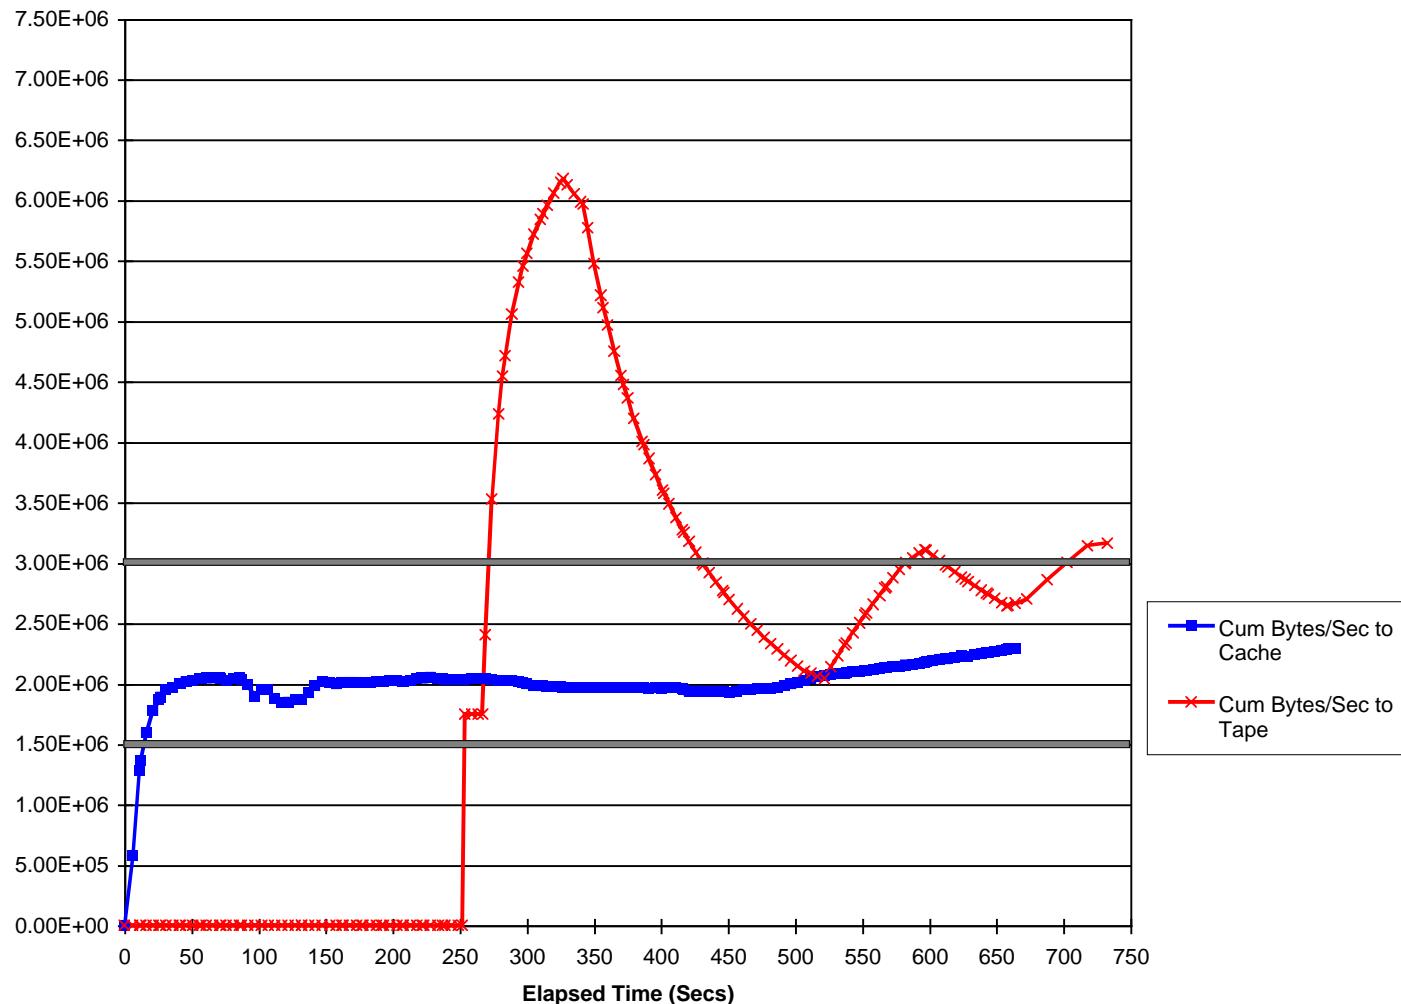
Copy from Ingest to AMASS Cache to AMASS Tape
No ftp Contention at Ingest Source



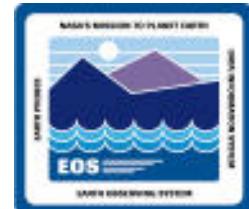
With ftp Contention—Rates



Copy from Ingest to AMASS Cache to AMASS Tape
With ftp Contention at Ingest Source



Next Steps: Source-to-Ingest-to-Archive Thread



Rerun test at GSFC

- RCC code to be installed after August Demo
- Expect improvement from better hardware and network
 - Ingest processor has more CPU (2 vs. 1 processors & faster clock rate) and more disk I/O bandwidth (2 vs. 1 controller)
 - Archive processor has more disk I/O bandwidth (8 vs. 4 controllers)

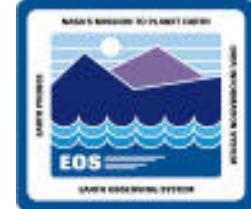
Next, add a run with multiple granules

- Expect improvement in total throughput from parallel threads

Later, rerun with simultaneous reads out of the tape archive

- Investigate impacts of contention for AMASS tape drives and/or AMASS disk cache

Ingest Test Conclusions



Rates

- Without ftp contention: **3.1 MB/sec**
- With ftp contention (@4 MB/sec) **2.3 MB/sec**
- Blended rate: **> 2.5 MB/sec**

For single-threaded ingest-only operation in mini-DAAC configuration, AMASS tape can write faster than network/AMASS cache can transfer

Current sustained rate of 2.3 MB/sec supports ingest of 198 GB of data per day

- MODIS at GSFC needs only 70 GB/day at launch
- Landsat at EDC needs only 139 GB/day
- MISR, MOPPIT, and CERES at LaRC need only 49 GB/day
- ASTER at EDC is ingested via D3 tape, following a different thread

Measured Source-to-Ingest-to-Archive rates

- Exceed the August Demo success criterion
- Will improve just by moving to real DAACs
- Represent reasonable progress